



Overview of Midrange Computing Resources at LBNL

Gary Jung
March 26, 2002

A Look at Existing Midrange Resources



Purpose of this exercise

- Illustrate the usage and demand for midrange computing cycles at LBNL
 - Lab usage of NERSC computing resources
 - Division-owned resources
- Show which platforms/architectures are used.
- Observe current trends

What is of interest?



NERSC Resources

- Find out number of projects using NERSC
- Look at allocation


Division-owned resources

- Systems larger than a desktop.
 - 4 or more processors/nodes
- Linux Clusters
- Distributed memory systems
- SMP systems

LBNL Use of Scientific Computing Resources


GFLOPS
Billions of Floating Point
Operations per Second
(Peak Rating)

5 TFLOPS



NERSC-3 IBM SP RS/6000 3328 processors
50 LBNL projects awarded 9,471,200 MPP hrs
through NERSC standard allocation process in FY02

500 GFLOPS




NERSC-LBNL T3E agreement (FY98-FY00)
12 LBNL projects awarded 150,000 MPP hrs FY00

**100 - 200
GFLOPS**


**Small to Medium SMP
systems**
UX8 General Timesharing,
DSD Diesel project, Joint
Genome Institute,
Computational Crystallography
Initiative, Berkeley Drosophila,
Yucca Mountain,
Computational Geophysics



**IBM SP RS/6000
16-40 processors**
Theresa Head-
Gordon Group,
Martin Head-
Gordon Group



**LBNL 80 node 160 processor
"alvarez" Linux Cluster**
2 LBNL projects



**8-36 node Linux Clusters
(High Performance Interconnect)**
Yucca Mountain, Computational
Geophysics, Future Technologies PC
Cluster and Babel Cluster



PDSF 207 node 390 processor Linux Cluster
14 LBNL High Energy Physics and Nuclear
Physics collaborations including STAR, E895,
E871, Amanda, ATLAS, Babar, SNO



**Small 10-32 node Linux Clusters
(Fast Ethernet Interconnect)**
Drosophila, SuperNova, Lighting Simulation,
Steve Brenner Research Group, Numerical
Algorithm Group, Sung Hou Kim Group, Joint
Genome Institute, Center for Functional Imaging,
NERSC Data Management



**Tightly Coupled
Architecture**

**Loosely Coupled
Architecture**

Findings - NERSC Usage



- Use of NERSC resources through the standard allocation process has been increasing.
 - 50 LBNL projects awarded 709,850 MPP hrs in FY00
 - 52 projects awarded 2,128,822 MPP hrs in FY01
 - 50 projects awarded 9,471,200 MPP hrs in FY02
- Use of NERSC resources through the LBNL T3E agreement increased each year during the FY98-FY00 period.
 - 13 LBNL projects awarded 86,000 MPP hrs in FY98
 - 15 LBNL projects awarded 100,000 MPP hrs in FY99
 - 12 LBNL projects awarded 150,000 MPP hrs in FY00
- NERSC PDSF used by many Physics and Nuclear Science projects including STAR, E871, E895, ATLAS, Amanda, Babar, SNO.

Findings - Division-owned Resources



Architecture	Division	Project	# Nodes or CPUS	Model	CPU	Interconnect
Small Linux clusters	Life Sciences	Center for Functional Imaging	17	Intel	Pentium P4	100BT
	Energy and Environment	Lighting Simulation	10	AMD	1.2Ghz Athlon	100BT
	Genome	Joint Genome Institute	16	AMD	2 x 1.2Ghz Athlon	100BT
	Life Sciences	Berkeley Drosophila Genome Project	32	Intel	2x 700Mhz PIII	100BT
	NERSC	Numerical Algorithm	12	AMD	1.3Ghz Athlon	100BT
	NERSC	Data Management	8	AMD	1 x 1800+ Athlon	100BT
	Physical Biosciences	Steve Brenner	80	Intel	2x1Ghz PIII	100BT
	Physical Biosciences	Sung Hou Kim	6	Intel	2x1Ghz PIII	100BT
	Physics	SuperNova	9	Intel	6 @ 300Mhz; 3 @ 600Mhz PII	100BT
Small Linux clusters with high performance interconnect	Earth Sciences	Yucca Mountain	20	Intel	2x1Ghz PIII	Gigabit
	Earth Sciences	Center for Computational Geophysics	8	Intel	2x1Ghz PIII	Myrinet & 100BT
	NERSC	Future Technologies PC Cluster	36	Intel	400Mhz PII	Myrinet & 100BT
	NERSC	Future Technologies Babel Cluster	12	Alpha	466Mhz Alpha	Myrinet & Gigabit
Distributed memory systems	Life Sciences	Theresa Head-Gordon	40	IBM SP RS/6000	375Mhz IBM Power3 CPUs	SP Switch & 100BT
	Chemical Sciences	Martin Head-Gordon	18	IBM Pseries	IBM Power3 CPUs	100BT
			16	IBM SP2	IBM Power3 CPUs	SP Switch
Architecture	Division	Project	# of Systems	Model	# and type of CPUs	
Small to medium SMP systems	Earth Sciences	Center for Computational Geophysics	1	Sun E4000	4 ea. 336Mhz Ultrasparc II CPUs	
		Yucca Mountain Project	1	Sun E4000	14 ea. 248 Mhz Ultrasparc II CPUs	
	ITSD	UX8 General Timesharing	1	Sun E4500	6 ea. 400Mhz Ultrasparc II CPUs	
	NERSC	Distributed Systems Dept Diesel project	2	Sun E4000	8 ea. 400Mhz Ultrasparc II CPUs	
			2	Sun E4500	4 ea. 400Mhz Ultrasparc II CPUs	
	Genome	Joint Genome Institute	1	Sun E6500	20 ea. 360Mhz Ultrasparc II CPUs	
			1	Sun E3000	8 ea. 400Mhz Ultrasparc II CPUs	
			1	Sun E450	4 ea. 400Mhz Ultrasparc II CPUs	
	Physical Biosciences	Computational Crystallography Initiative	2	Compaq ES40	4x833Mhz Alpha	
			1	Compaq DS10	466Mhz Alpha	
			1	Compaq XP900	466Mhz Alpha	
			1	Compaq DS20E	2x667Mhz Alpha	

Observations



- Increased usage of NERSC resources through the standard allocation process and PDSF.
- Purchases of large SMP systems have declined as evidenced by survey. Most existing SMP systems surveyed are utilizing older generation processors.
- Recently, the trend has been for Divisions to purchase and operate their own Linux clusters. Most of these are loosely coupled systems utilizing fast ethernet as the interconnect.
- Fewer purchases of parallel clusters utilizing a high speed interconnect such as Myrinet because of increased costs and complexity.

Example of a Division-owned Resource



Berkeley Drosophila Genome Project

Usage: Genomic sequence annotation using BLAST.

Hardware: LINUX Cluster by LINUX Networkx. 20 nodes, 2 x 700Mhz Intel PIII CPUs, 512MB/node

Interconnect: 100BaseT interconnect

Scheduler: Batch process using PBS

- Procurement
 - 2 months to determine RFP
 - 3 weeks to get quotes from 4 vendors and determine vendor
 - 2 months to get funding approval (outside agency)
- Setup:
 - 3 weeks in semi operations.
 - 2 months to get in shape for users
 - 0.75 FTE effort
- Purchase Costs:
 - \$65K (\$95K - \$30K rebate because of academic/showcase.)
 - \$3K vendor installation fee
- Added another 12 nodes Aug 2001 to bring total nodes to 32. \$32K purchase including vendor installation